

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	16928	(database or (data adj base)) near3 (monitor\$4 or manag\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:07
S2	783486	valid or secure	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 11:59
S3	1999017	valid or secur\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 12:00
S4	4210	(reject\$4 or den\$4) adj2 (list or table)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 12:01
S5	1090	S1 same S3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:11
S6	0	S5 same S4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:11
S7	22	S5 and S4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:11
S8	380284	registry	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:01
S9	2	"6038563".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:02

EAST Search History

S10	1	S8 and S9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:11
S11	1	S10 and S5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:11
S12	571563	ID	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:03
S13	559773	key	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:03
S14	0	S12 and S11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:11
S15	1	S13 and S11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S16	4442	key adj value	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:08
S17	1553	process adj ID	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:09
S18	0	registry adj key	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:09
S19	28	leaking adj data	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:09

EAST Search History

S20	204655	database	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:09
S21	351	registry adj key	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/01/07 13:10
S22	0	S19 and S20 and S21	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S23	3	S19 and S20	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S24	0	(Starek near Phillip).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:09
S25	3	(Murdock near Carlos).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:10
S26	37	(Friedman near George).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:10
S27	24563	(database or (data adj base)) near3 (monitor\$4 or manag\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S28	2307710	valid or secur\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S29	1698	S27 same S28	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12

EAST Search History

S30	5328	(reject\$4 or den\$4) adj2 (list or table)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S31	1698	S27 same S28	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S32	0	S31 same S30	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S33	28	S31 and S30	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S34	438667	registry	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S35	2	"6038563".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S36	1	S34 and S35	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S37	1	S34 and S35	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S38	1	S37 and S31	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S39	1	S37 and S31	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12

EAST Search History

S40	845044	ID	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S41	0	S40 and S39	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S42	685860	key	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S43	1	S42 and S39	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S44	38	leaking adj data	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S45	656923	database	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S46	536	registry adj key	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S47	0	S44 and S45 and S46	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:12
S48	7	S44 and S45	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:13
S49	38	S24 or S25 or S26	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:13

EAST Search History

S50	1	S48 and S49	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/18 11:13
-----	---	-------------	---	----	-----	------------------



Welcome United States Patent and Trademark Office

[Search Session History](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Fri, 31 Mar 2006, 3:03:04 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Search Query Display

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries

- | | |
|---------------------|--|
| #1 | ((registry key)<in>metadata) |
| #2 | (khandle<IN>metadata) |
| #3 | (database management<IN>metadata) |
| #4 | (security clearance<IN>metadata) |
| #5 | (security threshold<IN>metadata) |
| #6 | (security clearance<IN>metadata) |
| #7 | (adjustable threshold<IN>metadata) |
| #8 | (rejection list<IN>metadata) |
| #9 | handle |
| #10 | ((registry key)<in>metadata) <AND> (handle) <AND> ((database management<IN>metadata)) |
| #11 | ((((registry key)<in>metadata)) <AND> (handle) <AND> ((database management<IN>metadata))) <AND> ((security threshold<IN>metadata)) |

Indexed by
 Inspec

[Help](#) [Contact Us](#) [Privac](#)

© Copyright 2006 IE


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[handle](#) [database](#) [registry](#) [value](#) [adjustable](#) [threshold](#)

Found 33 of 173,942

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 33

Result page: [1](#) [2](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐

1 [GPGPU: general purpose computation on graphics hardware](#)



David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

 August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**
Publisher: ACM Press

 Full text available: [pdf\(63.03 MB\)](#) Additional Information: [full citation](#), [abstract](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

2 [Fast detection of communication patterns in distributed executions](#)



Thomas Kunz, Michiel F. H. Seuren

 November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**
Publisher: IBM Press

 Full text available: [pdf\(4.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

3 [Grapevine: an exercise in distributed computing](#)



Andrew D. Birrell, Roy Levin, Michael D. Schroeder, Roger M. Needham

 April 1982 **Communications of the ACM**, Volume 25 Issue 4

Publisher: ACM Press

 Full text available: [pdf\(1.71 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Grapevine is a multicomputer system on the Xerox research internet. It provides facilities for the delivery of digital messages such as computer mail; for naming people, machines,

and services; for authenticating people and machines; and for locating services on the internet. This paper has two goals: to describe the system itself and to serve as a case study of a real application of distributed computing. Part I describes the set of services provided by Grapevine and how its data and funct ...

4 A method for automatic rule derivation to support semantic query optimization

 Michael Siegel, Edward Sciore, Sharon Salveter
December 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 4

Publisher: ACM Press

Full text available:  pdf(2.73 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


The use of inference rules to support intelligent data processing is an increasingly important tool in many areas of computer science. In database systems, rules are used in semantic query optimization as a method for reducing query processing costs. The savings is dependent on the ability of experts to supply a set of useful rules and the ability of the optimizer to quickly find the appropriate transformations generated by these rules. Unfortunately, the most useful rules are not always th ...

Keywords: integrity constraint, learning, transformation heuristic

5 Watermarking relational data: framework, algorithms and analysis

Rakesh Agrawal, Peter J. Haas, Jerry Kiernan
August 2003 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 12 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(223.17 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)


Abstract. We enunciate the need for watermarking database relations to deter data piracy, identify the characteristics of relational data that pose unique challenges for watermarking, and delineate desirable properties of a watermarking system for relational data. We then present an effective watermarking technique geared for relational data. This technique ensures that some bit positions of some of the attributes of some of the tuples contain specific values. The specific bit locations and value ...

Keywords: Database, Information hiding, Steganography, Watermarking

6 Experience with Grapevine: the growth of a distributed system

 Michael D. Schroeder, Andrew D. Birrell, Roger M. Needham
February 1984 **ACM Transactions on Computer Systems (TOCS)**, Volume 2 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.54 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

Keywords: Grapevine

7 Experiences with selecting search engines using metasearch

 Daniel Dreilinger, Adele E. Howe
July 1997 **ACM Transactions on Information Systems (TOIS)**, Volume 15 Issue 3

Publisher: ACM Press

Full text available:  pdf(428.65 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)

[terms](#), [review](#)

Search engines are among the most useful and high-profile resources on the Internet. The problem of finding information on the Internet has been replaced with the problem of knowing where search engines are, what they are designed to retrieve, and how to use them. This article describes and evaluates SavvySearch, a metasearch engine designed to intelligently select and interface with multiple remote search engines. The primary metasearch issue examined is the importance of carefully selecti ...

Keywords: WWW, information retrieval, machine learning, search engine

8 [Visualizing geospatial data](#)



Theresa Marie Rhyne, Alan MacEachern, Theresa-Marie Rhyne

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**

Publisher: ACM Press

Full text available: [pdf\(13.99 MB\)](#) Additional Information: [full citation](#), [abstract](#)

This course reviews concepts and highlights new directions in GeoVisualization. We review four levels of integrating geospatial data and geographic information systems (GIS) with scientific and information visualization (VIS) methods. These include:• Rudimentary: minimal data sharing between the GIS and Vis systems• Operational: consistency of geospatial data• Functional: transparent communication between the GIS and Vis systems• Merged: one comprehensive toolkit environmentW ...

9 [Mobile data management: Middleware support for reconciling client updates and data transcoding](#)



Thomas Phan, George Zorpas, Rajive Bagrodia

June 2004 **Proceedings of the 2nd international conference on Mobile systems, applications, and services MobiSys '04**

Publisher: ACM Press

Full text available: [pdf\(4.80 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In mobile Internet applications, data can be transcoded, updated, and transferred across heterogenous clients. The problem then arises where updates made in the context of an initial transcoding results in content too stringently transcoded for subsequent clients, thereby causing loss of semantic value. We solve this problem by suggesting that the updates themselves can be transformed so that they can be applied directly to the original data instead of to the transcoded data; this approach allow ...

Keywords: client updates, data management, middleware, mobile computing, reconciliation, transcoding

10 [Stateful distributed interposition](#)



John Reumann, Kang G. Shin

February 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 1

Publisher: ACM Press

Full text available: [pdf\(833.84 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex servic ...

Keywords: Distributed computing, component services, distributed context, multitiered services, operating systems, server consolidation

11 Manageability, availability, and performance in porcupine: a highly scalable, cluster-based mail service



Yasushi Saito, Brian N. Bershad, Henry M. Levy

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Publisher: ACM Press

Full text available: [pdf\(2.52 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the motivation, design and performance of Porcupine, a scalable mail server. The goal of Porcupine is to provide a highly available and scalable electronic mail service using a large cluster of commodity PCs. We designed Porcupine to be easy to manage by emphasizing dynamic load balancing, automatic configuration, and graceful degradation in the presence of failures. Key to the system's manageability, availability, and performance is that sessions, data, and underlying ...

Keywords: cluster, distributed systems, email, group membership protocol, load balancing, replication

12 Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining



Robbert Van Renesse, Kenneth P. Birman, Werner Vogels

May 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: [pdf\(341.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Scalable management and self-organizational capabilities are emerging as central requirements for a generation of large-scale, highly dynamic, distributed applications. We have developed an entirely new distributed information management system called Astrolabe. Astrolabe collects large-scale system state, permitting rapid updates and providing on-the-fly attribute aggregation. This latter capability permits an application to locate a resource, and also offers a scalable way to track sys ...

Keywords: Aggregation, epidemic protocols, failure detection, gossip, membership, publish-subscribe, scalability

13 Research sessions: data integration: Adapting to source properties in processing data integration queries



Zachary G. Ives, Alon Y. Halevy, Daniel S. Weld

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: [pdf\(197.27 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)


An effective query optimizer finds a query plan that exploits the characteristics of the source data. In data integration, little is known in advance about sources' properties, which necessitates the use of *adaptive* query processing techniques to adjust query processing on-the-fly. Prior work in adaptive query processing has focused on compensating for delays and adjusting for mis-estimated cardinality or selectivity values. In this paper, we present a generalized architecture for adaptiv ...

14 Data integrity: Web application security assessment by fault injection and behavior monitoring

Yao-Wen Huang, Shih-Kun Huang, Tsung-Po Lin, Chung-Hung Tsai

May 2003 **Proceedings of the 12th international conference on World Wide Web**

Publisher: ACM Press

Full text available:  [pdf\(4.53 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As a large and complex application platform, the World Wide Web is capable of delivering a broad range of sophisticated applications. However, many Web applications go through rapid development phases with extremely short turnaround time, making it difficult to eliminate vulnerabilities. Here we analyze the design of Web application security assessment mechanisms in order to identify poor coding practices that render Web applications vulnerable to attacks such as SQL injection and cross-site scr ...

Keywords: black-box testing, complete crawling, fault injection, security assessment, web application testing

15 B2B e-commerce and enterprise integration: The development and evaluation of exception handling mechanisms for order fulfillment process based on BPEL4WS

Fu-ren Lin, Hsiang-chin Chang

August 2005 **Proceedings of the 7th international conference on Electronic commerce ICEC '05**

Publisher: ACM Press

Full text available:  [pdf\(426.00 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The advance of Internet technology fosters the order fulfillment process in a supply chain across heterogeneous information systems. In order to monitor states between partners in the process, this study develops exception handling mechanisms based on Web service stack. After detecting exceptions, the first step is to make use of the fault and compensation handlers provided by the BPEL4WS (Business Process Execution Language for Web Services) specification to roll back planned or even executed b ...


Keywords: BPEL4WS, Web service, exception handling, resource management, supply chain management

16 Cellular disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(287.05 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource management, scalable multiprocessors, virtual machines

17 Application servers, enterprise computing, and software engineering: Extending a J2EE™ server with dynamic and flexible resource management

Mick Jordan, Grzegorz Czajkowski, Kirill Kouklinski, Glenn Skinner

October 2004 **Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware Middleware '04**

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(407.32 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The Java™ 2 Platform, Enterprise Edition (J2EE™) is the standard platform for hosting enterprise applications written in the Java programming language. A single J2EE server can support multiple applications much like a traditional operating system, but performance levels can be difficult to control, due to the absence of resource management facilities in the Java platform. The Resource Management (RM) interface addresses this problem by providing a flexible and extensible framework f ...

18 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP '99**, Volume 33 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.93 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a syste ...

19 An integrated platform for reliable multicast support in the regional mobile-IP environment

Hassan Omar, Tarek Saadawi, Myung Lee

April 2002 **ACM SIGMOBILE Mobile Computing and Communications Review**, Volume 6 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(167.80 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Supporting reliable delivery of multicast datagrams, in IP networks, may necessitate the introduction of new elements and features. Further, considerable additional signaling may be required to support this service. Providing a platform that efficiently supports IP multicast delivery, in an environment where the multicast group members frequently change their locations, is a challenge for systems supporting mobility. In this paper, we describe a platform that allows the application of an interna ...

20 Service discovery in agent-based pervasive computing environments

Olga Ratsimor, Dipanjan Chakraborty, Anupam Joshi, Timothy Finin, Yelena Yesha

December 2004 **Mobile Networks and Applications**, Volume 9 Issue 6

Publisher: Kluwer Academic Publishers

Full text available:  [pdf\(443.01 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Directory based service discovery mechanisms are unsuitable for ad-hoc m-commerce environments. Working towards finding an alternate mechanism, we developed Allia: a peer-to-peer caching based and policy-driven agent-service discovery framework that

facilitates cross-platform service discovery in ad-hoc environments. Our approach achieves a high degree of flexibility in adapting itself to changes in ad-hoc environments and is devoid of common problems associated with structured compound forma ...

Keywords: advertising, agents, caching, device preferences, device profiles, mobile service discovery

Results 1 - 20 of 33

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)